arXiv:1605.01746v1 [cs.NE] 5 May 2016

Biobjective Performance Assessment with the COCO Platform

Dimo Brockhoff[1], Tea Tušar[1], Dejan Tušar[1], Tobias Wagner[2],
Nikolaus Hansen[3,4], Anne Auger[3,4]

[1]Inria, research centre Lille, [2]Technical University Dortmund,
[3]Inria, research centre Saclay, [4]Université Paris-Saclay, LRI

Abstract

This document details the rationales behind assessing the performance of numerical black-box optimizers on multi-objective problems within the COCO platform and in particular on the biobjective test suite `bbob-biobj`. The evaluation is based on a hypervolume of all non-dominated solutions in the archive of candidate solutions and measures the runtime until the hypervolume value succeeds prescribed target values.

Contents

1 Introduction **2**
1.1 Definitions and Terminology 2

2 Performance Assessment with a Quality Indicator **3**
2.1 Definition of the Quality Indicator 3
2.2 Rationales Behind the Performance Measure 4
2.3 Specificities and Properties 5

3 Definition of Target Values **6**
3.1 Target Precision Values 6
3.2 The Reference Hypervolume Indicator Value 7

4 Instances and Generalization Experiment **7**

5 Data Storage and Future Recalculations of Indicator Values **7**

1 Introduction

The performance assessment of (numerical) optimization algorithms with the COCO platform *[HAN2016co]* is invariably based on the measurement of the *runtime*[1] until a *quality indicator* reaches a predefined *target value*. On each problem instance, several target values are defined and for each target value a runtime is measured (or no runtime value is available if the indicator does not reach the target value) *[HAN2016perf]*. In the single-objective, noise-free case, the assessed quality indicator is, at each given time step, the function value of the best solution the algorithm has obtained (evaluated or recommended, see *[HAN2016ex]*) before or at this time step.

In the bi- and multi-objective case, e.g. on the biobjective `bbob-biobj` test suite *[TUS2016]*, the assessed quality indicator at the given time step is a hypervolume indicator computed from *all* solutions obtained (evaluated or recommended) before or at this time step.

1.1 Definitions and Terminology

In this section, we introduce the definitions of some basic terms and concepts.

function instance, problem In the case of the bi-objective performance assessment within COCO, a problem is a 5-tuple of

- a *parameterized function* $f_\theta : \mathbb{R}^n \to \mathbb{R}^2$, mapping the decision variables of a solution $x \in \mathbb{R}^n$ to its objective vector $f_\theta(x) = (f_\alpha(x), f_\beta(x))$ with $f_\alpha : \mathbb{R}^n \mapsto \mathbb{R}$ and $f_\beta : \mathbb{R}^n \mapsto \mathbb{R}$ being parameterized (single-objective) functions themselves
- its concrete parameter value $\theta \in \Theta$ determining the so-called *function instance* i,
- the *problem dimension* n,
- an underlying quality indicator I, mapping a set of solutions to its quality, and
- a *target value* I_{target} of the underlying quality indicator, see below for details.

We call a problem *solved* by an optimization algorithm if the algorithm reaches a quality indicator value at least as good as the associated target value. The number of function evaluations needed to surpass the target value for the first time is COCO's central performance measure. *[HAN2016co]* In case a single quality indicator is used for all problems in a benchmark suite, we can drop the quality indicator and refer to a problem as a quadruple $f_\theta, \theta, n, I_{\text{target}}$. Note that typically more than one problem for a *function instance* of (f_θ, θ, n) is defined by choosing more than one target value.

Pareto set*, *Pareto front*, and *Pareto dominance For a function instance, i.e., a function $f_\theta = (f_\alpha, f_\beta)$ with given parameter value θ and dimension n, the Pareto set is the set of all (Pareto-optimal) solutions for which no solutions in the search space $\mathbb{R}^n$ exist that have either an improved f_α or an improved f_β value while the other value is at least as good (or in other

[1] Time is considered to be *number of function evaluations* and, consequently, runtime is measured in number of function evaluations.

words, a *Pareto-optimal* solution in the Pareto set has no other solution that *dominates* it). The image of the Pareto set in the *objective space* is called the Pareto front. We generalize the standard Pareto dominance relation to sets by saying solution set $A = \{a_1, \ldots, a_{|A|}\}$ dominates solution set $B = \{b_1, \ldots, b_{|B|}\}$ if and only if for all $b_i \in B$ there is at least one solution a_j that dominates it.

ideal point The ideal point (in objective space) is defined as the vector in objective space that contains the optimal function value for each objective *independently*, i.e. for the above concrete function instance, the ideal point is given by $z_{\text{ideal}} = (\inf_{x \in \mathbb{R}^n} f_\alpha(x), \inf_{x \in \mathbb{R}^n} f_\beta(x))$.

nadir point The nadir point (in objective space) consists in each objective of the worst value obtained by any Pareto-optimal solution. More precisely, if $\mathcal{PO}$ denotes the Pareto set, the nadir point satisfies $z_{\text{nadir}} = (\sup_{x \in \mathcal{PO}} f_\alpha(x), \sup_{x \in \mathcal{PO}} f_\beta(x))$.

archive An external archive or simply an archive is the set of non-dominated solutions, obtained over an algorithm run. At each point t in time (that is after t function evaluations), we consider the set of all mutually non-dominating solutions that have been evaluated so far. We denote the archive after t function evaluations as A_t and use it to define the performance of the algorithm in terms of a (quality) indicator function $A_t \to \mathbb{R}$ that might depend on a problem's underlying parameterized function and its dimension and instance.

2 Performance Assessment with a Quality Indicator

For measuring the runtime on a given problem, we consider a quality indicator which is to be optimized (minimized). In the noiseless single-objective case, the quality indicator is the best so-far observed objective function value (recommendations can replace previous observations). In the case of the `bbob-biobj` test suite, the quality indicator is based on the hypervolume indicator of the *archive* A_t.

2.1 Definition of the Quality Indicator

The indicator $I_{\text{HV}}^{\text{COCO}}$ to be mininized is either the negative hypervolume indicator of the archive with the nadir point as reference point or the distance to the region of interest $[z_{\text{ideal}}, z_{\text{nadir}}]$ after a normalization of the objective space[2]:

$$I_{\text{HV}}^{\text{COCO}} = \begin{cases} -\text{HV}(A_t, [z_{\text{ideal}}, z_{\text{nadir}}]) & \text{if } A_t \text{ dominates } \{z_{\text{nadir}}\} \\ dist(A_t, [z_{\text{ideal}}, z_{\text{nadir}}]) & \text{otherwise} \end{cases} .$$

where

$$\text{HV}(A_t, z_{\text{ideal}}, z_{\text{nadir}}) = \text{VOL}\left(\bigcup_{a \in A_t} \left[\frac{f_\alpha(a) - z_{\text{ideal},\alpha}}{z_{\text{nadir},\alpha} - z_{\text{ideal},\alpha}}, 1 \right] \times \left[\frac{f_\beta(a) - z_{\text{ideal},\beta}}{z_{\text{nadir},\beta} - z_{\text{ideal},\beta}}, 1 \right] \right)$$

[2] We conduct an affine transformation of both objective function values such that the ideal point $z_{\text{ideal}} = (z_{\text{ideal},\alpha}, z_{\text{ideal},\beta})$ is mapped to $(0, 0)$ and the nadir point $z_{\text{nadir}} = (z_{\text{nadir},\alpha}, z_{\text{nadir},\beta})$ is mapped to $(1, 1)$.

is the (normalized) hypervolume of archive A_t with respect to the nadir point $(z_{\text{nadir},\alpha}, z_{\text{nadir},\beta})$ as reference point and where (with division understood to be element-wise, Hadamard division),

$$dist(A_t, [z_{\text{ideal}}, z_{\text{nadir}}]) = \inf_{a \in A_t, z \in [z_{\text{ideal}}, z_{\text{nadir}}]} \left\| \frac{f(a) - z}{z_{\text{nadir}} - z_{\text{ideal}}} \right\|$$

is the smallest (normalized) Euclidean distance between a solution in the archive and the region of interest, see also the figures below for an illustration.

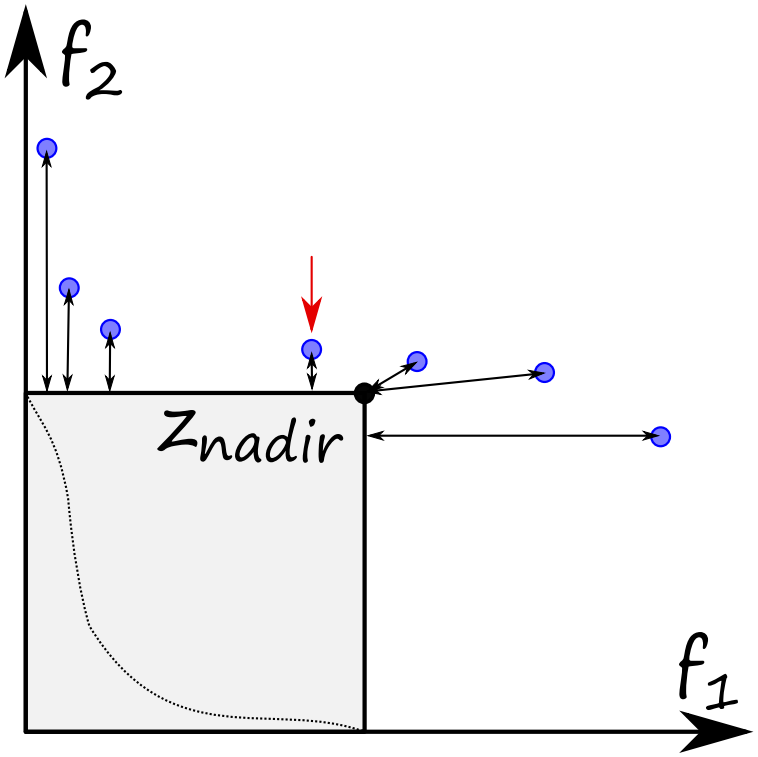

Fig. 1: Illustration of Coco's quality indicator (to be minimized) in the (normalized) bi-objective case if no solution of the archive (blue filled circles) dominates the nadir point (black filled circle), i.e., the shortest distance of an archive member to the region of interest (ROI), delimited by the nadir point. Here, it is the fourth point from the left (indicated by the red arrow) that defines the smallest distance.

2.2 Rationales Behind the Performance Measure

Why using an archive? We believe using an archive to keep all non-dominated solutions is relevant in practice in bi-objective real-world applications, in particular when function evaluations are expensive. Using an external archive for the performance assessment has the additional advantage that no populuation size needs to be prescribed and algorithms with different or even changing population sizes can be easily compared.

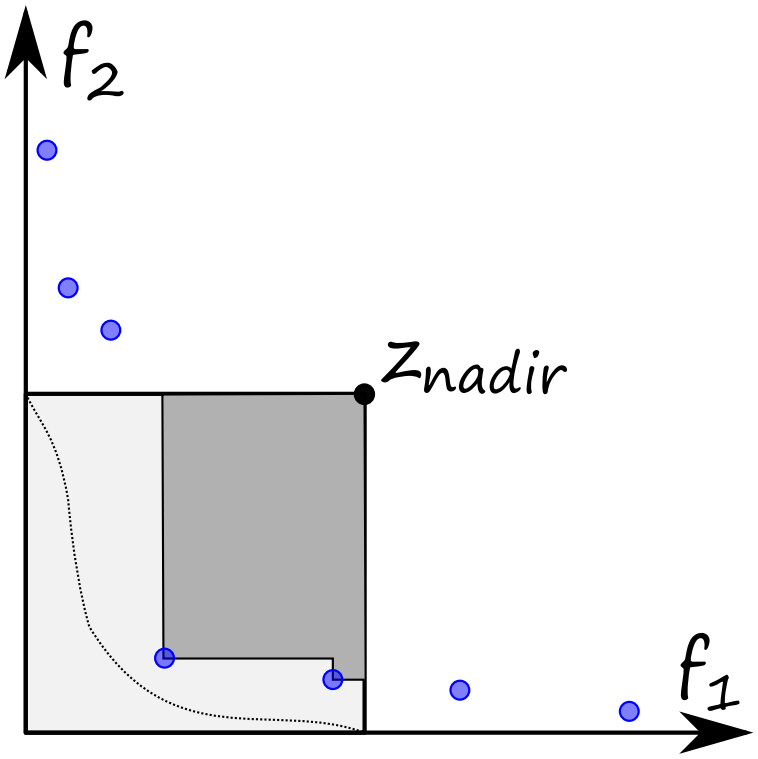

Fig. 2: Illustration of Coco's quality indicator (to be minimized) in the bi-objective case if the nadir point (black filled circle) is dominated by at least one solution in the archive (blue filled circles). The indicator is the negative hypervolume of the archive with the nadir point as reference point.

Why hypervolume? Although, in principle, other quality indicators can be used in replacement of the hypervolume, the monotonicity of the hypervolume is a strong theoretical argument for using it in the performance assessment: the hypervolume indicator value of the archive improves if and only if a new non-dominated solution is generated *[ZIT2003]*.

2.3 Specificities and Properties

In summary, the proposed `bbob-biobj` performance criterion has the following specificities:

- Algorithm performance is measured via runtime until the quality of the archive of non-dominated solutions found so far surpasses a target value.
- To compute the quality indicator, the objective space is normalized. The region of interest (ROI) $[z_{\text{ideal}}, z_{\text{nadir}}]$, defined by the ideal and nadir point, is mapped to $[0, 1]^2$.
- If the nadir point is dominated by at least one point in the archive, the quality is computed as the negative hypervolume of the archive using the nadir point as hypervolume reference point.
- If the nadir point is not dominated by the archive, the quality equals the distance of the

archive to the ROI.

This implies that:

- the quality indicator value of an archive that contains the nadir point as non-dominated point is 0.
- the quality indicator value is bounded from below by -1, which is the quality of an archive that contains the ideal point, and
- because the quality of an archive is used as performance criterion, no population size has to be prescribed to the algorithm. In particular, steady-state and generational algorithms can be compared directly as well as algorithms with varying population size and algorithms which carry along their external archive themselves.

3 Definition of Target Values

For each problem instance of the benchmark suite, consisting of a parameterized function, its dimension and its instance parameter θ_i, a set of quality indicator target values is chosen, eventually used to measure algorithm runtime to reach each of these targets. The target values are based on a target precision ΔI and a *reference hypervolume indicator value*, I_i^{ref}, which is an approximation of the $I_{\mathrm{HV}}^{\mathrm{COCO}}$ indicator value of the Pareto set.

3.1 Target Precision Values

All target indicator values are computed in the form of $I_i^{\mathrm{ref}} + \Delta I$ from the instance dependent reference value I_i^{ref} and a target precision value ΔI. For the `bbob-biobj` test suite, 58 target precisions ΔI are chosen, identical for all problem instances, as

$$\Delta I \in \{\underbrace{-10^{-4}, -10^{-4.2}, \ldots, -10^{-4.8}, -10^{-5}}_{\text{six negative target precision values}}, 0, 10^{-5}, 10^{-4.9}, 10^{-4.8}, \ldots, 10^{-0.1}, 10^{0}\} \ .$$

Negative target precisions are used because the reference indicator value, as defined in the next section, can be surpassed by an optimization algorithm.[3] The runtimes to reach these target values are presented as empirical cumulative distribution function, ECDF *[HAN2016perf]*. Runtimes to reach specific target precisions are presented as well. It is not uncommon however that the quality indicator value of the algorithm never surpasses some of these target values, which leads to missing runtime measurements.

[3] In comparison, the reference value in the single-objective case is the f-value of the known global optimum and, consequently, the target precision values have been strictly positive *[HAN2016perf]*.

3.2 The Reference Hypervolume Indicator Value

Unlike the single-objective bbob test suite *[HAN2009fun]*, the biobjective bbob-biobj test suite does not provide analytic expressions of its optima. Except for f_1, the Pareto set and the Pareto front are unknown.

Instead of the unknown hypervolume of the true Pareto set, we use the hypervolume of an approximation of the Pareto set as reference hypervolume indicator value I_i^{ref}.[4] To obtain the approximation, several multi-objective optimization algorithms have been run and all non-dominated solutions over all runs have been recorded.[5] The hypervolume indicator value of the obtained set of non-dominated solutions, also called *non-dominated reference set*, separately obtained for each problem instance in the benchmark suite, is then used as the reference hypervolume indicator value.

4 Instances and Generalization Experiment

The standard procedure for an experiment on a benchmark suite, like the bbob-biobj suite, prescribes to run the algorithm of choice once on each problem of the suite *[HAN2016ex]*. For the bbob-biobj suite, the postprocessing part of COCO displays currently by default only 5 out of the 10 instances from each function-dimension pair.

5 Data Storage and Future Recalculations of Indicator Values

Having a good approximation of the Pareto set/Pareto front is crucial in assessing algorithm performance with the above suggested performance criterion. In order to allow the reference sets to approximate the Pareto set/Pareto front better and better over time, the COCO platform records every non-dominated solution over the algorithm run. Algorithm data sets, submitted through the COCO platform's web page, can therefore be used to improve the quality of the reference set by adding all solutions to the reference set which are currently non-dominated to it.

Recording every new non-dominated solution within every algorithm run also allows to recover the algorithm runs after the experiment and to recalculate the corresponding hypervolume difference values if the reference set changes in the future. In order to be able to distinguish between different collections of reference sets that might have been used during the actual benchmarking experiment and the production of the graphical output, COCO writes the absolute hypervolume reference

[4] Using the quality indicator value of the *true* Pareto set might not be desirable, because the set contains an infinite number of solutions, which is neither a possible nor a desirable goal to aspire in practice.

[5] Amongst others, we run versions of NSGA-II *[DEB2002]* via Matlab's gamultiobj function, SMS-EMOA *[BEU2007]*, MOEA/D *[ZHA2007]*, RM-MEDA *[ZHA2008]*, and MO-CMA-ES *[VOS2010]*, together with simple uniform RANDOMSEARCH and the single-objective CMA-ES *[HAN2001]* on scalarized problems (i.e. weighted sum) to create first approximations of the bi-objective problems' Pareto sets.

values together with the performance data during the benchmarking experiment and displays a version number in the plots generated that allows to retrieve the used reference values from the Github repository of COCO.

Acknowledgements

This work was supported by the grant ANR-12-MONU-0009 (NumBBO) of the French National Research Agency.

The authors would like to thank Thanh-Do Tran for his contributions and assistance with the preliminary code of the bi-objective setting and for providing us with his extensive experimental data. We also thank Tobias Glasmachers, Oswin Krause, and Ilya Loshchilov for their bug reports, feature requests, code testing, and many valuable discussions. Special thanks go to Olaf Mersmann for the inital rewriting of the COCO platform without which the bi-objective extension of COCO would not have happened.

References

[BEU2007] N. Beume, B. Naujoks, and M. Emmerich (2007). SMS-EMOA: Multiobjective selection based on dominated hypervolume. *European Journal of Operational Research*, 181(3), pp. 1653-1669.

[DEB2002] K. Deb, A. Pratap, S. Agarwal, and T. A. M. T. Meyarivan (2002). A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2), pp. 182-197.

[HAN2001] N. Hansen and A. Ostermeier (2001). Completely derandomized self-adaptation in evolution strategies. *Evolutionary computation*, 9(2), pp. 159-195.

[HAN2016perf] N. Hansen, A. Auger, D. Brockhoff, D. Tušar, T. Tušar (2016). COCO: Performance Assessment, *ArXiv e-prints*, arXiv:160x.xxxxx.

[HAN2016co] N. Hansen, A. Auger, O. Mersmann, T. Tušar, D. Brockhoff (2016). COCO: A Platform for Comparing Continuous Optimizers in a Black-Box Setting, *ArXiv e-prints*, arXiv:1603.08785.

[HAN2009fun] N. Hansen, S. Finck, R. Ros, and A. Auger (2009). Real-parameter black-box optimization benchmarking 2009: Noiseless functions definitions. Technical Report RR-6829, Inria, updated February 2010.

[HAN2016ex] N. Hansen, T. Tušar, A. Auger, D. Brockhoff, O. Mersmann (2016). COCO: The Experimental Procedure, *ArXiv e-prints*, arXiv:1603.08776.

[TUS2016] T. Tušar, D. Brockhoff, N. Hansen, A. Auger (2016). COCO: The Bi-objective Black Box Optimization Benchmarking (bbob-biobj) Test Suite, *ArXiv e-prints*, arXiv:1604.00359.

[VOS2010] T. Voß, N. Hansen, and C. Igel (2010). Improved step size adaptation for the MO-CMA-ES. In *Genetic and Evolutionary Computation Conference (GECCO 2010)*, pp. 487-494. ACM.

[ZHA2007] Q. Zhang, and H. Li (2007). MOEA/D: A multiobjective evolutionary algorithm based on decomposition. *IEEE Transactions on Evolutionary Computation*, 11(6), pp. 712-731.

[ZHA2008] Q. Zhang, A. Zhou, and Y. Jin (2008). RM-MEDA: A regularity model-based multiobjective estimation of distribution algorithm. *IEEE Transactions on Evolutionary Computation*, 12(1), pp. 41-63.

[ZIT2003] E. Zitzler, L. Thiele, M. Laumanns, C. M. Fonseca, and V. Grunert da Fonseca (2003). Performance Assessment of Multiobjective Optimizers: An Analysis and Review. *IEEE Transactions on Evolutionary Computation*, 7(2), pp. 117-132.